

Present and Future Computing Requirements for m526: Computational Resources for the Nanomaterials Theory Institute at the Center for Nanophase Materials Sciences

Paul R. C. Kent

CNMS, ORNL

NERSC BES Requirements for 2017

October 8-9, 2013

Gaithersburg, MD

1. Project Description


PI, Co-PIs: Kent (PI)+CNMS NTI Staff (Co-PIs)

- Predictive calculations into the behavior & properties of nanosystems: graphene, catalysts, molecular electronics, nanomaterials synthesis...
- Supports CNMS user projects as well as CNMS theme science (internal). Leads to significant diversity in scientific questions, tools/applications use, and computational workload.
- General trend: more realism to match experimental conditions.
- Simulations are primarily atomistic using either quantum mechanics or classically-based potentials.

Theory of high-rate lithium intercalation pseudocapacitance in Nb_2O_5

Experiment: V. Augustyn et al. Nature Materials (June 2013); shows high rate intercalation, high energy density of Nb_2O_5 , apparently capacitive.

Theory: Need to determine lithium locations, and energy barriers for transport, all as a function of lithium concentration, for two crystal phases.



Requires: DFT calculations for crystal structure, construction of a cluster expansion (~100 DFT calculations), nudged elastic band for energy barriers, longer DFT-based molecular dynamics to check transport (few long calculations). All on a few nodes on hopper/edison, hours to a few days total runtime.

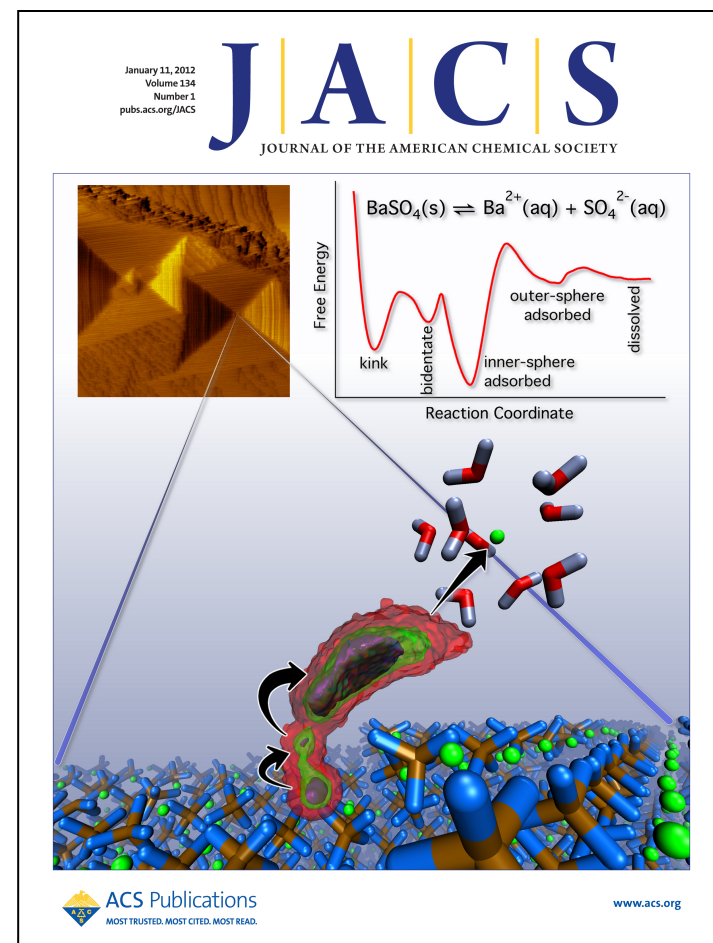
Issues: Most calculations on simpler variant of experimental structure to reduce cost.

Publication: A. Lubimtsev et al., J. Mat. Chem. A. (Accepted, October 2013)

Discovering crystal growth mechanisms

Theory: A grand challenge-level problem is to determine the mechanisms that influence (nano)materials growth and dissolution. Can not be accessed through direct molecular dynamics because the processes are too slow. Accelerated MD, metadynamics and related “rare event” theories enable these mechanisms to be accessed.

Requires: Many long molecular dynamics runs (many ns), here using LAMMPS with custom force-field.
Issues: (i) More complex mechanisms &/or materials require longer runs. (ii) Rare event theories – by exploring more of configuration space – find problems with underlying potentials! Techniques demand overall increased accuracy of potentials and ideally use strong experimental calibrations.



Stack et al. JACS (2012).
Isosurfaces of Free Energy for
Ba ion removal from BaSO_4 step

1. Project Description (continued)

- Scientific goals for 2017 include:
 - I. (Routine) Discovery and exploration of properties of not-yet synthesized materials, *using simulations of sufficient sophistication that we can have confidence in predictions*
 - Implies routine benchmarking & validations of predictions using the technique “one better” in accuracy.
 - For the most part, these techniques exist already, but are increasingly expensive.
 - E.g. Down-selecting which functionalized graphene analogs should be synthesized from many possibilities.

1. Project Description (continued)

- II. A more thorough understanding and exploration of already identified materials, chemical systems
- Easier access to longer time or length scales will facilitate integration with neutron scattering experiments
 - Use of rare event theories (accelerated MD) enables access to processes not accessible by direct MD.

1. Project Description (continued)

III. Overall increased confidence in predictions through use of improved methods

- More accurate, formerly specialist methods are becoming commoditized
- Requires larger allocations: greater accuracy has greater cost.
- Community standards are increasing. Requested by referees.
- E.g. Instead of qualitative DFT band gaps and trends, now using ~quantitative GW routinely, but N^6 scaling instead of N^3 (N =no of atoms); hybrid DFT for thermodynamics (N^4 vs N^3); self-consistent van der Waals/dispersion (2-4x)

2. Computational Strategies

- Approximately 50% of our allocation is DFT, 25% classical MD, 25% other methods (quantum chemistry, quantum Monte Carlo, phase field modeling...)
- We approach these problems computationally at a high level primarily by adopting either classical or quantum-mechanics based dynamics:
 - Classical MD primarily uses LAMMPS
 - Quantum mechanics is primarily computed using density functional theory (DFT), normally using one of the many variants of the plane-wave pseudopotential approximation. E.g. Quantum Espresso, VASP, ABINIT.
[Expect more local basis set methods (CP2K etc.) in future]
- !! Important to note !! : None of these codes should be considered as monolithic. E.g. The DFT codes all implement hybrid functionals, which have very different costs+performance to conventional DFT.

2. Computational Strategies (cont)

- These codes are characterized by:
 - LAMMPS: domain decomposed classical MD (many options, many force-fields)
 - QE, VASP, ABINIT: MPI distributed FFTs, dense linear algebra, non-linear optimization
- Our biggest computational challenges are ...
 - Cost. Individual investigations are increasingly expensive.
- Our parallel scaling is limited by ...
 - Difficulty of hiding communications (Classical MD, local DFT)
 - Immature implementations (post local DFT methods)
- In 2017: Expect to use refined/improved implementations, but not fundamentally different. (Could discuss quantum chemistry here; DFT might not improve fast enough).

3. Current HPC Usage (see slide notes)

- Machines currently using: edison, hopper, carver
- Hours used in 2012-2013: ~10M @ NERSC
- Typical parallel concurrency and run time, number of runs per year
 - Tens of thousands of runs, increasingly via automation.
 - Estimate: mean concurrency ~256 cores, 95 % of runs ≤ 10000 cores (long tail distribution)
 - Codes checkpoint
- Data read/written per run
 - 1-10s Gigabytes
- Memory used per (node | core | globally)
 - 1-64 Gbytes/node

4. HPC Requirements for 2017

(Key point is to directly link NERSC requirements to science goals)

- Compute hours needed (in units of Hopper hours):
 - Each 2017 goal readily implies a 10x+ increase in required compute hours:
 - VALIDATION requires the use of more expensive techniques, e.g. quantum chemistry to validate density functional predictions.
 - RARE EVENT theories – typically implemented as long molecular dynamics runs with biases dynamics. May require more accurate (costly) potentials.
 - TRANSITION to HIGHER ACCURACY – hybrid DFT and GW readily 10x more expensive, even for small systems (scaling with system size is higher, prefactor is higher)
- Changes to parallel concurrency, run time, number of runs per year
 - More expensive (worse formal scaling) techniques also scale better due to increased work. E.g. A reasonable hybrid DFT implementation can utilize thousands of cores, even for small system sizes (see Bylaska's implementation in NWCHEM, recent CPMD, QBox)
- Changes to data read/written
 - In proportion with system size, # systems computed, timescale.
- Changes to memory needed per (core | node | globally)
 - Per node memory ~64GBytes ~acceptable provided interconnect is good.

5. Strategies for New Architectures

- Our strategy for running on new many-core technologies (GPUs or MIC) is:
 - We are “developer limited”: carefully focus our limited development time
 - Adopt open source applications, as much as possible. Leverage world-wide GPU/MIC effort.
 - Attempt to encourage users to use open source
- We have prepared for many core by: thoroughly understanding algorithms and bottlenecks of main simulation codes; choosing codes with good existing threaded &/or GPU implementations; polling users for usability & performance comments on simulation codes
- We are already planning to:
 - Improve the usability & implementation of QE
- To be successful on many-core systems we will need help with:
 - Development. Coordination of development. (To a good approximation, there are no developers, only users.)
 - Profiling, performance monitoring (improving, but still awkward with real data)

5. Summary

- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
 - We are compute limited: “Whatever provides the greatest increase in user compute hours in a not-too hostile form”
 - Any new architecture/machine should be both energy *and* developer efficient
- NERSC generally refreshes systems to provide on average a 2X performance increase every year. What significant scientific progress could you achieve over the next 5 years with access to 32X your current NERSC allocation?
 - (Many possible examples due to broad scope of our project)
 - E.g. A modern, atomistic theory of simple crystal growth/dissolution, relevant to materials growth, scale formation in oil pipelines.... Build on prior work on BaSO₄ growth (Stack JACS 2012) and ongoing user project work for CaCO₃. Current problem: metadynamics finds/highlights issues with the underlying potential... need additional validation, refit &/or use of more expensive+accurate potentials (ReaxFF? DFT?)